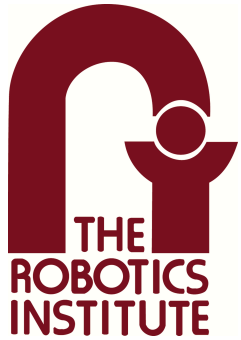

Carnegie Mellon University
The Robotics Institute, School of Computer Science

MRSD Individual Lab Report

ILR04

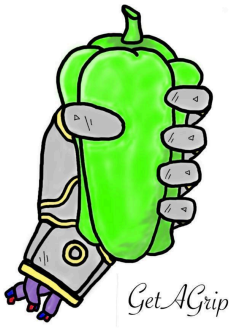

Team D: GetAGrip.AI

Teammates: Solomon Fenton, Jiyoon Park, Shri Ishwaryaa S V, Sridevi Kaza, Alec Trela

Author: Shri Ishwaryaa S V

Submission Date: Mar. 23, 2023

1 Individual Progress:

1.1 MRSD Project:

The following tasks have been completed by me between 3/2/23 and 3/23/23:

1) Trained U-Net with a ResNet34 backbone and a weighted loss function for peduncle segmentation
2) Found a dataset for plant stem segmentation
3) Set up the depth camera to connect with my laptop and display RGB-D data
4) Obtained the point of interaction with the peduncle from the peduncle mask

Some additional progress made by Jiyoon in the perception subsystem are listed below

1) Trained YOLOv8 for peduncle segmentation
2) Set up SegNet for plant stem segmentation
3) Developing the GUI for our system with Alec

In the previous ILR, I mentioned that I will improve the accuracy of the peduncle segmentation and work on writing code in the perception pipeline. I trained U-Net again with a ResNet34 backbone and used a weighted loss function which weighs an error in the peduncle class more than the background class. This improved the performance which was presented in the Preliminary Design Review. Jiyoon was also training YOLOv8 parallelly for peduncle segmentation. YOLOv8 gives more reliable results on our testbed as shown in Fig. 1. In case we need to make our model robust, we can combine both the U-Net and YOLOv8 to obtain final results.

Figure 1: YOLOv8's peduncle segmentation results on our testbed setup

The next task on our list was to identify the main stem of the plant and pass its location as a constraint to the motion planning subsystem. To perform stem segmentation, we

needed a good dataset. Fortunately, we found a dataset consisting of 10,000 images with the annotation in the form of a black-and-white mask. I am working on training YOLOv8 on this data. The masks need to be converted to the YOLOv8 annotation format. This has been a roadblock for me because no software performed this conversion readily. Hence, I have written code to do the conversion and am currently debugging it.

One other task that I was able to do was connect the RealSense D435i camera to my laptop by completing the necessary installations mentioned on the official GitHub page. An image of the RealSense viewer displaying the RGB-D data is shown in Fig. 2. Future steps include establishing communication between the camera and ROS.

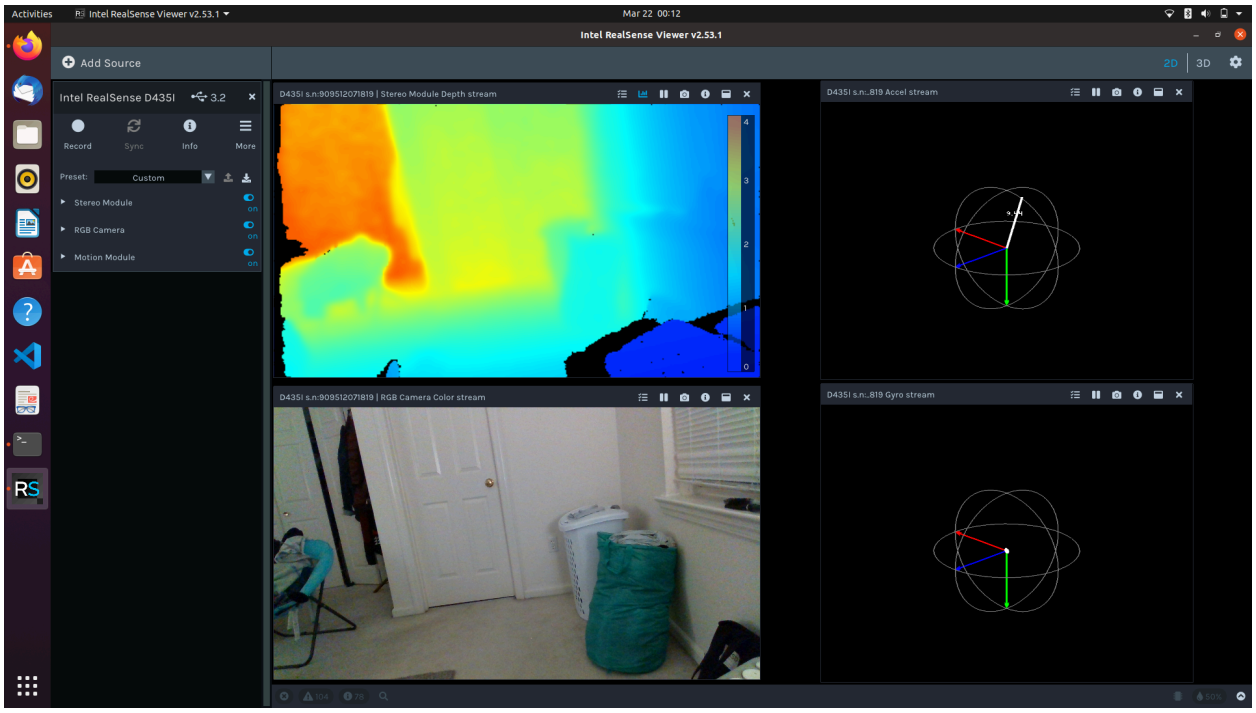

Figure 2: Intel RealSense viewer for the D435i camera

Finally, I wrote the pipeline code for calculating the point of interaction with the peduncle from the peduncle mask obtained from YOLOv8. Firstly, a morphological closing is performed to remove any dark spots in the masks. Then, scipy's medial_axis function is used to obtain the medial axis of the peduncle mask. A parabolic curve is fit to the medial axis data. Now, we choose a point on the curve at a threshold distance from the top of the pepper. During the presentation, I had a constant threshold value in terms of pixels. However, the code has been modified so that the threshold is 50% of the total peduncle length. These results are displayed in Fig. 3.

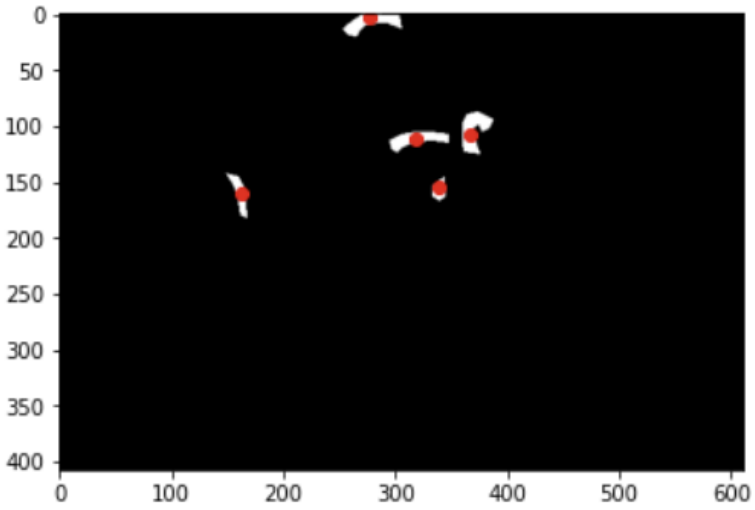

Figure 3: POI with a threshold of 50% of total peduncle length

2 Challenges:

2.1 MRSD Project:

A major challenge I faced was finding a dataset for plant stem segmentation. Luckily, Jiyoon found a Kaggle page with a dataset and code for training SegNet on it. She focused on training SegNet and faced some issues which I won't be going over in this report. I was trying to use the data to train YOLOv8. As I mentioned already, converting the black and white masks to YOLOv8 format has been a challenge. This is because we are uploading the data to Roboflow which performs some data processing for us. I have written code that outputs the annotation as a text file that YOLOv8 requires. When I upload these text files to Roboflow, it is not recognising them as annotations. Therefore, I am currently trying different popular annotation formats to see which one is compatible with Roboflow.

I will upload them in a format that Roboflow accepts and convert them directly to YOLOv8's format in Roboflow itself.

3 Teamwork:

3.1 MRSD Project:

Alec Trela

Alec has set up a personal 3D printer in his house so that we can make steady progress in the end-effector subsystem without encountering any failed prints. Using this, he has reduced the size of the cutter. In addition, he attempted to tune the PID position control of the motor. He also executed a cutting fidelity test to validate our cutter design. The results are promising.

Solomon Fenton

One item that Solomon completed is achieving the ideal TPU surface pad texture. With this, he could grip a peduncle strongly even when the gripper was shaken vigorously. This marks the completion of designing the gripping mechanism. He also performed the end-effector error-handling test to determine which point on the cutter is best for cutting.

Sridevi Kaza

Sri spent a lot of time trying to connect Sawyer to ROS. After many meetings with Francisco, a student in the FRC, and MRSD seniors, she was able to do it. She is currently troubleshooting moving the physical arm with ROS and MoveIt. While debugging this issue, she was able to simulate the arm moving to a desired position with a constraint on its final orientation.

Jiyoon Park

Jiyoon finished training YOLOv8 for peduncle segmentation. She found SegNet for plant stem segmentation and tested the network. She collaborated with Alec to develop the GUI for our system. Furthermore, as the only person on the team with extensive experience with physical manipulators, she is helping Sri debug all the motion planning subsystem issues. She also created a Docker image for the arm. Lastly, she has written some pipeline code for pepper detection.

Team

Since we are near the integration of some subsystems, the team developed a ROS architecture for the entire system. We also had more collaborations between Sri, Alec, and Jiyoon. To allow everyone to participate in other subsystems, we have assigned small tasks to each one of us outside our primary subsystems. For example, Alec is working on the GUI, Sri and Solomon will work on writing some of the perception code, and Jiyoon and I will be making the mount for the end-effector.

4 Plans:

4.1 MRSD Project:

Tasks that I wish to achieve before the next progress review are:

1) Train YOLOv8 for stem segmentation
2) Complete the entire pipeline with Jiyoon

Jiyoon has defined classes named pepper and peduncle, and written some of the pipeline code. I have finished the code for obtaining a POI. Now, functions such as prioritizing the order in which peppers need to be picked must be completed. The perception subsystem aims to complete the entire pipeline before the next progress review. We hope to be in the middle of system integration by then.